

Appendix B: Additional background information and figures to the pipeline.

The orbital fat deformation visualization pipeline, see figure 1, was developed using the C++, OpenGL, GLSL and CUDA programming tools, and binaries are available for Linux and Windows. It consisted of three phases: (1) data acquisition and processing, (2) specification of a ROI, and (3) visualization of sliding and deformation. Each of these phases in turn consisted of a number of steps, together comprising a hierarchical framework that we used to investigate fat deformation in the orbit.

1. Data Acquisition and Processing

In the first phase of the pipeline, we obtained digital representations of physical quantities present in our research domain. For 3D fat deformation visualization from static MRI data, this included a sequence of MRI data volumes containing the research domain, and time-dependent deformation data reflecting its changing structure.

1.1. MRI acquisition.

During the first phase, MRI data of the deforming anatomy is acquired by making acquisitions in multiple directions of gaze as outlined in the methods section of the main article.

1.2. Deformation calculation

We applied to this sequence two different deformation calculation methods, and visualized the results (figure 1, panel 1b). Deformation data expresses the incremental structural changes in data volumes as vector fields. This deformation data can be used to visualize patches of soft tissue changing their position and shape over time. Optical Flow methods (Beauchemin & Barron, 1995) result in a sequence of deformation velocity fields, where each vector estimates the direction and degree at which a local patch of tissue is deforming at a specific instant. Registration methods (Maintz & Viergever, 1998) (Dawant, 2002), on the other hand, result in deformation-offset fields.

The used optical flow method was 3D Lucas and Kanade Optical Flow (LK3D) (Lucas & Kanade, 1981) which results in a sequence of 15 deformation data volumes of the same dimensions and resolution as the input data. The registration method consisted of B-Splines Mutual Information Non-Rigid Registration (B-Splines) (Klein, et al., 2007), resulting in a sequence of 14 deformation offset volumes, also of the same dimensions and resolution as the input data. The registration is based on a mutual information metric and the deformation is parameterized by B-splines.

When comparing LK3D and B-Splines in time-dependent deformation, the difference between the two becomes apparent. Tracking a single point through the data domain over a single unit of time using velocity fields requires stepwise integration using time interpolated deformation velocity data, whereas offset fields describe instant displacement over unit time, i.e. from one acquisition volume to the next. Deformation offset data can be used as deformation velocity data, provided that between acquisition steps the velocity data is not interpolated. However, this causes sudden velocity changes at acquisition step boundaries, which is detrimental to fluent visualization. Interestingly, in some cases B-Splines proved to be superior (figures 2, 3, 4 and 5), whilst in others LK3D better represents sliding and

deformation (figures 6 and 9). The comparison between the two is described in more detail in Appendix B.

The extrapolation, however, of deformation patterns based on the MRI images (figure 1, panel 1a and 1b) was affected by several factors. First, the sequence of data volumes should match the result of 3D snapshots taken at regular intervals from a smoothly deforming volume. As the acquisition of a single MRI volume took several minutes, we assumed that the deformation of orbital fat depended only on the instantaneous eyeball rotation angle, and not on the angular velocity of the rotating eyeball. Also, presence of noise or lack of spatial resolution limited detection of small anatomical features, and consequently their contribution to the deformation data.

As noise in the MRI data and characteristics of the deformation calculation method cause inaccuracies in the deformation data. We had no way of a priori calculating the reliability of the deformation data, although Abramoff and Viergever used a metric derived from the LK3D algorithm to estimate uncertainty (Abramoff & Viergever, 2002). We assessed the correctness of an individual deformation field by using it to deform its associated data volume over unit time, and comparing it to the next data volume in the sequence. This led to a scalar uncertainty field, (figure 1, panel 1c) which we used as instantaneous visual feedback on deformation reliability during visualization.

2. Specification of a Region of Interest

ROI specification comprises the selection of a sub-region of a data volume that will receive visualization focus. It is necessary in dense 3D visualization to prevent visuals in irrelevant regions from occluding the currently more interesting parts of a 3D volume. It also prevents overwhelming the user with visual detail by limiting representation only to those parts of the volume the user deems interesting. We used ROIs to demarcate a sub-volume of fat, the deformation of which we would visualize.

We assumed that the method illustrated in the pipeline (figure 1, panels 2a through 2c) would allow quick selection of virtually arbitrary shapes in the data domain by (1) drawing an arbitrary plane through the dataset on a surface showing prominent anatomical features, (2) fine-tuning its placement and (3) extruding the outline of an intended ROI (Schaafsma, et al., 2010).

2.1 Drawing a plane

Before a plane is drawn, only a Reference Surface model, or R-surface is visible. This R-surface serves as a simple, yet sufficient representation of the bony orbit, eyeballs and optic nerves and thereby aids in the 3D orientation in the 3D orbit and adequate plane positioning. In figure B1, left image, one can see clearly the dark gray R-surface showing both orbits and surrounding tissue. Inside of both the eyeballs are visible. Superimposed and in light gray is part of the right orbit. In the bottom-left corners of the images, the orientation and appearance of the *dummy* representation of both eyeballs and optic nerves match the current view and angle of eyeball rotation. In the central image, an inferior axial view of the right orbit, as also indicated by the *dummy* model in the bottom-left corner, is presented. The eyeball is clearly visible as the spherical shape to the right, and connected to it is the optic nerve, which will be moving downwards in this view as eyeball rotation progresses. In the right image, a vertical anterior

view of the right eyeball and a small part of the optic nerve are shown. Without the model this view is hard to interpret.

Inside the R-surface model, a plane is drawn which serves as base for the ROI. MRI data is projected on to the plane, making the plane an arbitrarily oriented 2D view through a 3D volume. We refer to this plane as the Interaction plane, or I-plane. Figure B2 shows, in a more detailed fashion, how the I-plane is drawn. As can be seen, the R-surface functions as a 3D orientation cue in which a red dotted line represents the drawing pattern by the user. As a result of this, an I-plane is formed. In the left image, the intersection of the spherical eyeball becomes approximately circular when placing an I-plane perpendicular to the viewing direction. As shown in the central image, an I-plane perpendicular to the optic nerve can be positioned by choosing an appropriate viewpoint (inset) and drawing a straight line on the R-surface. In the right image, an axial I-plane, approximately the plane of rotation of the right eyeball, containing the optic nerve, is positioned. The inset on the left shows the original drawing view.

2.2 Fine-tuning using a L-Widget

Drawing the intersection between I-plane and R-surface alone is not enough to provide the accuracy and flexibility needed for positioning arbitrary anatomically relevant planes. The intersection may be inaccurately estimated or indicated, or the R-surface might not include the features needed to draw the intended I-plane. Translation or predictable rotation of the I-plane can assist in both these matters, which we facilitate with an in-view graphical interaction element, or widget (Markoupolos, Rowson, & Johnson, 1997) (Wang, Sajeew, & Inchaiwong, 2006). We have developed the L-widget for this purpose (Figure B3). With its base glued to the I-plane it serves as a handle to translate and rotate the I-plane (figure 1, panel 2b).

An I-plane is created with the L-widget already in place. Drawing on the plane repositions the L-widget. This gives the user control over the L-widget rotation axes, while translation remains unaffected. Grabbing and dragging the spherical handles manipulates the L-widget, and hence the I-plane. As a result the L-widget allows fast and intuitive positioning of meaningful I-planes in the dataset, and assists in exploring the raw MRI data in a fashion similar to Multi-Planar Reconstruction by dragging the I-plane through the dataset (Figures B4 and B5).

2.3 Extruding the ROI

After accurately positioning the I-plane, the outline for the eventual ROI is drawn on the I-plane. The user extrudes this outline perpendicular to the I-plane using an extrusion handle that appears when the user completes the outline (figure 1, panel 2c).

This concept is schematically illustrated in figure B6. In figure B7, some more complex, yet clinically more relevant, ROIs are formed. On the left image an I-plane was positioned perpendicular to the optic nerve, and a circular shape, drawn on the I-plane, was extruded to capture part of the optic nerve. The central image, however, shows a ROI just behind the eyeball, parallel to the axial eyeball rotation plane. In the right image, the L-widget was used to position the I-plane on a branching structure; a blood vessel (**V**). Part of this blood vessel was outlined and slightly extruded to create a ROI that bears no direct relation to the R-surface

As expected, the aforementioned methods resulted in a quick and accurate selection of a ROI in the data domain. Although the intersection-based I-plane specification was quite intuitive due to its relation with the available anatomical features, additional interaction was required for the fine-tuning of the I-plane position. The developed L-Widget served well at achieving this goal. Since the extrusion of the ROI resulted from free-hand drawing on the I-plane, we were able to specify arbitrarily shaped ROIs. Therefore, a wide range of anatomical features could be investigated. More importantly using a ROI during visualization significantly limited the visual details to which the user was exposed, thereby providing a clear and ordered representation of the section deemed interesting.

3. Used ROIs in the figures

3.1 Sliding and deformation of orbital fat near the eyeball (main article sections 2.2.1 and 3.2.1; figure 2)

To create a useful ROI for our purpose, we navigated towards a superior view of the eyeball, and created a proper I-plane by dragging a straight line perpendicular to the circular edge of the eyeball. A side view of the deforming tissue inside the ROI emphasizes deformation in the direction of rotation of the eyeball.

3.2 Sliding and deformation of orbital fat around the optic nerve (main article sections 2.2.2 and 3.2.2; figures 3, 4 and 5)

To investigate the fat's behavior in front of the optic nerve, we positioned an I-plane coinciding with the center axis of the optic nerve, perpendicular to the movement direction of the optic nerve. On this I-plane, we drew a rectangular shape near the eyeball and extruded it towards the medial rectus muscle.

For the purpose of visualizing the patterns behind the optic nerve we created two different ROIs. The first ROI was created by placing an I-plane coinciding with the optic nerve axis, normal to the direction of motion of the optic nerve. On this I-plane, approximately rectangular outlines were drawn and extruded upwards. Grid spacing was adjusted to distinguish between a random configuration of the markers and a stack configuration. The other ROI was created differently. Here, we placed an I-plane approximately tangent to the eyeball at the point where the optic nerve attaches to it, and drew the outline in the shape of the marker's mass.

We also applied our visualization tool to investigate the motion of orbital fat to the sides of the optic nerve. The marker sheets in this figure were created as ROIs. First, an I-plane next to the optic nerve was positioned, depending on whether anterior or posterior marker sheets were needed. On this I-plane, outlines that stretched significantly more along the optic nerve than perpendicular to it were drawn. Specifying a large value for the grid spacing parameter corresponding to the second principal component of the outline drawing pattern, ensured a single sheet of markers after extrusion.

3.3 Displacement of orbital fat between eyeball and medial- and lateral rectus muscles during abduction (main article sections 2.2.3 and 3.2.3; figure 6)

For this experiment, an approximately axial I-plane through the optic nerve and medial and lateral rectus muscles was created. To clear our view on the MRI data, we temporarily removed the R-surface,

after which the I-plane clearly showed the area where eyeball and eye muscle meet. Then, an outline was drawn on the I-plane, and after bringing back the R-surface, was extruded to a proper thickness. After the ROI was created, we switched our view from to within the eyeball for better observation of the deformation patterns.

4. Marker-Based Deformation Visualization

4.1 Marker properties

To visualize the shape and position of a region of orbital fat, we used a marker-based method. These markers are in fact little dots that initially fill the ROI (figure 1, panel 3a). Without further qualification, markers are simply a mass of dots without any relevant structure. By adding stretching according to their instantaneous degree of deformation, patterns became visible. The droplet shape, as opposed to simple stretching according to the deformation field tangent, disambiguated between opposing directions of motion (B8d). Although this last aspect was hardly necessary during animation, it did provide valuable detail when dealing with images or stills. Marker density and size were adjusted to balance ROI size and deformation pattern scale. Too small markers can become indistinguishable and move as a single texture, while too large markers make the deformation appear erratic, and harder to identify patterns in. In general, the spatial scale of the sought deformation patterns determined the proper marker granularity scale.

Marker granularity and the uncertainty field further qualified the deformation visualization. Increased marker uncertainty degrades marker appearance to a gray-colored splat; marker (*d*) has minimal uncertainty, while the motion for marker (*f*) is highly uncertain, and marker (*e*) is somewhere in between. The result of the chosen uncertainty visualization was that markers in areas of high uncertainty or low confidence became indistinguishable, as exemplified by the left-most part of the ROI in the bottom-left figure B9; they formed an opaque gray patch that camouflages marker identity and motion. As expected (Hauser, 2006), this hid those regions from user attention, and automatically drew user focus to regions of higher confidence

4.2 Seeding of the markers

Seeding refers to the process of filling the ROI with the markers. Three basic seeding patterns were available: Randomly seeding pattern (default), Regular Grid pattern and Marker Sheets (Figure B10). Markers were seeded according to user-controlled spacing parameters. During animation, the user had interactive control over the seeding density rather than the number of seeded markers directly. This allowed consistent visualization results when different-sized ROIs were used throughout an exploration session. Adjusting density caused a new set of markers to be seeded inside the original ROI.

A randomly seeded ROI (left image), best conveyed the deformation of the tissue region as a whole. However, the small relative changes are unlikely to stand out, especially in a still image. On the other hand, in an approximately box-shaped ROI seeded with a regular grid of markers, small changes are noticed due to the rigid grid structure (center image). The image on the right shows a hemispherical ROI, partitioned in five densely and randomly seeded marker sheets. Using such a seeding pattern,

emphasized deformation in the sheet stacking direction, but small changes in the plane of a sheet are not well noticeable.

Conclusions:

Our method allows, without relying on automated 3D MRI volume extraction, selection of arbitrarily oriented anatomically relevant regions inside a dataset. The method we demonstrated extends common techniques such as parameterized geometric shapes (cubes, boxes and cylinders) and techniques based on the three principal anatomical planes. Also, the appearance of individual markers and the way in which they are distributed into the ROI help to understand detailed aspects of the deformation patterns present in the anatomy under investigation (video 2)